On the Impact of IP Option Processing - Part 2

Mattia Rossi University of Innsbruck (AT)
Institute of Computer Science
csac3227@uibk.ac.at

Michael Welzl University of Innsbruck (AT)
Institute of Computer Science
michael.welzl@uibk.ac.at

ABSTRACT

The first technical report about the impact of IP-Option processing [2], has shown that packets with options are processed slightly slower than packets without options even though this difference is not really significant. The data were gathered by pinging about 27000 hosts two times: once in August 2002 and once in July 2003 with the same list of hosts. Due to a lot of timeouts and/or other network errors only a seventh of the pinged hosts were useable for the statistics, which still provide enough relevant information.

The aim of this new technical report is to show the differences between a current measurement (february to april 2004) and the other two earlier measurements. This new measurement should give more precise information about the IP-Option processing by using a completely new, larger and more current hostlist.

1. INTRODUCTION

The interesting results gained by the measurements in 2002 and 2003 [2], which used the same hostlist of **27689** hosts, prompted us to try a new measurement with a more current and larger hostlist, generated by the web-crawler *"larbin"* [1]. This extremely fast crawler generated a hostlist of **138416** hosts in a few hours. The ping process started in February 2004 and was ended in April 2004. **46326** of all hosts on the hostlist were pinged. These are almost twice as many as in the first two measurements, while the amount of answering hosts has more than tripled to **18050** hosts. The *extping* was used like in the last measurements, with 100 pings per host and ping type. The execution line is the following:

```
ext_ping -n 70179 -p 100 -i hostlist.dat>output.dat
```

The program was planned to ping only 70179 hosts, because otherwise the ping process would have lasted too long, but it was interrupted earlier due to time constraints. The program generated the typical extping data files: *hostlist.dat, result.dat, router.dat, error.dat.* Afterwards the statistic tool and the summary program filtered the necessary information shown in this report.

2. GENERATING THE HOSTLIST

2.1 The *larbin* webcrawler

The hostlist used in the measurement of 2004 was generated with the larbin webcrawler[1]. The larbin webcrawler is a highly customizable and powerful webcrawler with lots of options and controllable by a web-interface, written by Sebastien Ailleret. Even if the larbin webcrawler is not really well documented, it's not too difficult to use. With a few adaptions it generated a hostlist with more than **200000** raw entries within a few hours.

Two files were changed accordingly: *larbin.conf* in the main directory and *useroutput.cc* in the *src/interf* directory. The *larbin.conf* is the main configuration file to set ports for the web-interface, the start page (`www.yahoo.com` in this case), the number of parallel searches and the filetypes to skip.

The *useroutput.cc* is the file containing the output procedure called by the main function. By default this procedure is empty. Few lines of C++ code change this, and the hosts of the perused webpages are stored in a file called *ergebnis.dat.* At this time a raw hostlist is generated.

2.2 Refining the hostlist

The raw hostlist unfortunately contains some problems. The hosts are listed in canonical form, which makes it difficult to ascertain that no two hosts share the same ip adress. The following commandlines changed the canonical names into the appropriate ip adress, deleted all double entries and sorted the list.

```
while read line; do host $line | grep address | cut
-d ' ' -f 4 >> ergebnis_raw.dat; done<ergebnis.dat
&& sort ergebnis_raw.dat | uniq > ergebnis_uniq.dat
```

This process delivered a hostlist of **138416** hosts, all of them reachable at this time.

3. THE MEASUREMENT 2004

The measurement 2004 started in February 2004 and went on until April 2004. During this period, **46326** hosts of the hostlist were already pinged, so we decided to stop the ping process completely and to continue with the statistics. The pings per hosts were still **100** alternating between NONE and NOP as in the measurements of 2002 and 2003. The interesting numbers are: **46326** pinged hosts, **18050** answering hosts, duration of the ping process: **2 months**. Like the measurement in 2003, the ping process took place in Innsbruck, Austria, on a PC with a 100 MBit connection to the internet, which luckily caused a low standard deviation in the RTTs, making the results more reliable.

Now for the results of this measurement. Figure 1 shows the average RTT per path length of the average NONE, NOP and NOP-A pings, while figure 2 shows the median RTT.

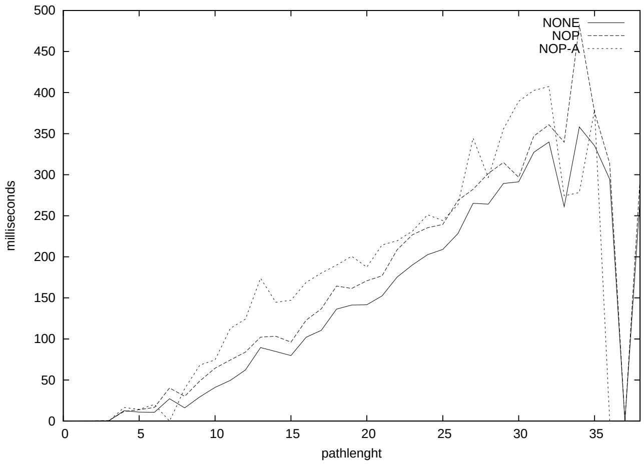

Figure 1: The average per-host-average results

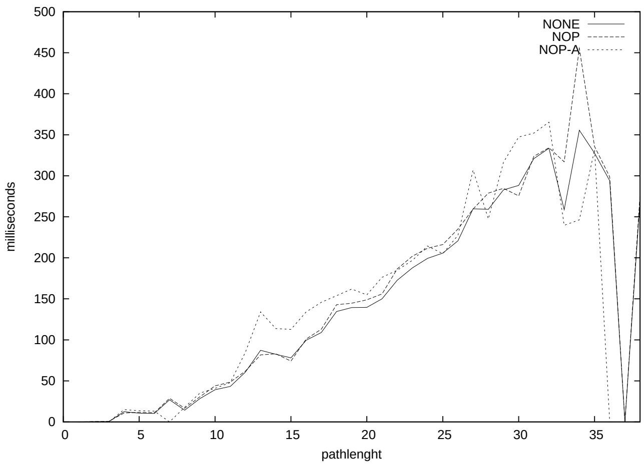

Figure 2: The average per-host-median results

Figure 1 clearly shows that the NONE Pings are processed faster than the NOP and NOP-A Pings. The NOP line never falls below the NONE line. The NOP-A line touches the zero line three times, because there are only few hosts with this pathlength, none of which answered with options. At pathlength 36 all lines touch the zero line, which means that there were no hosts with this pathlength. At one pathlength the highly oscillating NOP-A line falls below the NOP line, which was probably caused by a different return path. This is why NOP-A pings were not considered in our statistics.

The NOP line in the median graph is difficult to see, because it almost overlaps with the NONE but at some pathlengths it does fall slightly below the NONE. Even in this graph, the NOP-A line falls below the NOP line and even below the NONE line. As it can be seen in figure 3, there was only one host answering with a NOP-A at this pathlength. Therefore, this value does not really influence the statistics, as none of the results at the higher pathlengths do. Furthermore, figure 3 shows that the values of interest are approximately those from pathlength 8 to 32. The median graph shows that NONE and NOP values are generally not oscillating, and have no real strange behaviour in this segment. This means that it was a really clean and reliable measurement.

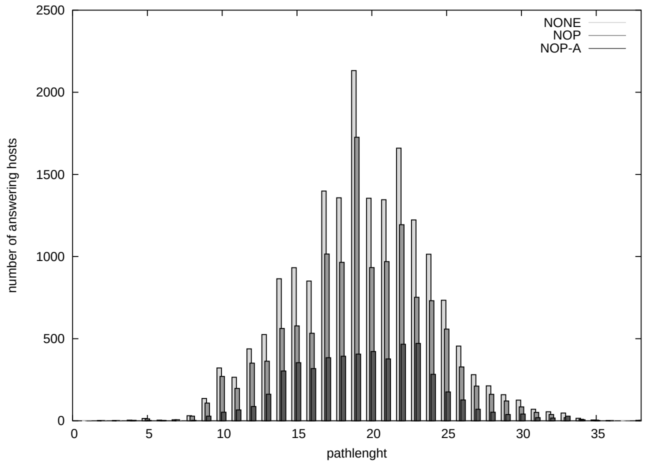

Figure 3: Answering hosts per path length

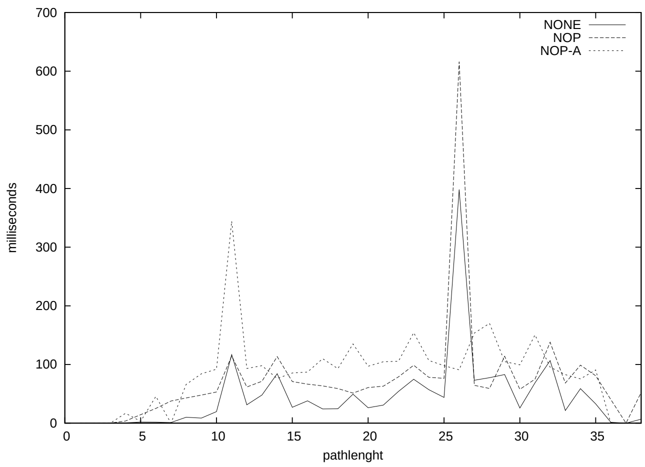

Figure 4: The average standard deviation

Figure 4 shows the standard deviation of the three ping types. The segment of main interest has two clear peaks. They are caused by some really slow routers or hosts. This should not influence the statistic, because hosts whose standard deviations are too high have approximately zero weight.

The interesting point is the persistent higher standard deviation of the NOP. This means that most routers behave similarly when processing NONE packets, which does not seem to be the case when processing NOP packets. The "`more error.dat|grep -c NOP`" command shows that **7545** routers or hosts still do not process NOP packets at all!

4. STATISTICS AND DIFFERENCES TO 2002 AND 2003

The statistics of this measurement confirm again the results already obtained by the measurements in 2002 and 2003. Figure 5 shows the weighted values.

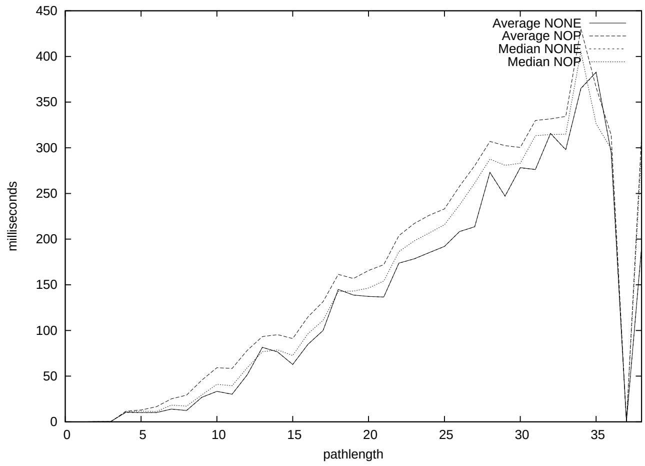

Figure 5: Averages and Medians per path length 2004

The most interesting aspect in this figure is the overlapping of the median NONE and the average NONE lines. Actually, the values are not exactly the same, but differ only by around 1 ms; since the values in this figure vary from 0 to 400, the differences cannot be seen. This characteristic shows that no greater interferences or errors were encountered during the ping process.

As for the measurements of 2002 and 2003 (fig. 6 and 7), the median NOP line is above the NONE line, and the average NOP line is above the median NOP line. The two lines are clearly separated from each other and from the NONE lines, whereas in 2002 and 2003 there were some oscillating values. The median NOP line falls below the NONE line only once, which is not surprising because the NONE line has a high peak at this pathlength as already described above. Figures 8 and 9 give a clearer impression of the coherence between the average NOP and the average NONE as well as the median NOP and the median NONE in the 2004 measurement.

Table 1 shows the final results of the statistics 2004, 2002 and 2003. The RTTs in this measurement are relatively low compared to the ones in the last measurements. Another interesting aspect is the small difference between the average and median NONE values, as already described above. The 2003 measurement shows this too. In 2002, the ping location did not have such a stable connection, so a higher difference was detected. The difference between the median NONE and the median NOP is **17.83 ms**,and between the averages it is **34.21 ms**. This is an average difference of **26.02 ms**.

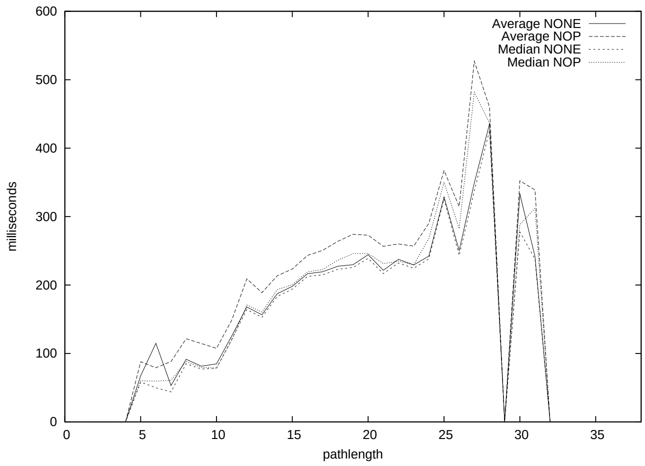

Figure 6: Averages and Medians per path length 2002

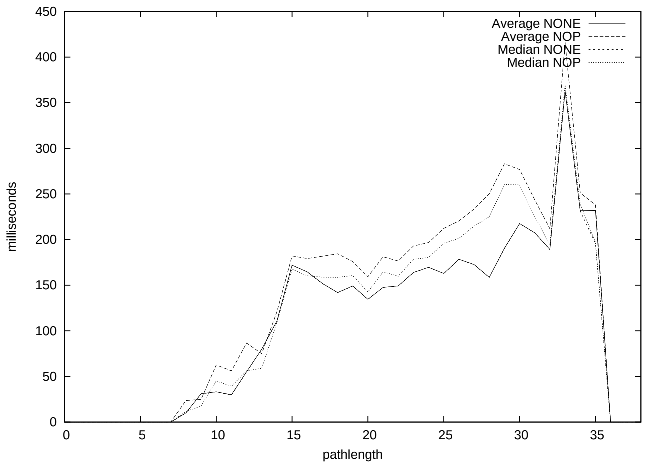

Figure 7: Averages and Medians per path length 2003

In 2002 the differences were **8.5 ms** (medians), **30.6 ms** (averages) and **19.55 ms**. In 2003, the values were **2.4 ms**,**19.95 ms** and **11.175 ms**. The value of the average difference of the 2004 measurement is the highest while the values of the NONE are the lowest, but the difference in percent is still only **26.5%** (**10%** in 2002, **7%** in 2003).

The host statistics show that the number of processed hosts is more than four times higher than in the measurements of 2002 and 2003, even if the number of pinged hosts is only doubled. Even the errors on NOP pings (routers/hosts that do not react on NOPs) are only doubled; these values are the result of using a more current hostlist. Interestingly, the paths of hosts that answered with a NOP on a NOP ping (NOP-A) had many routers, which indicates that the NOP-A pings chose relatively longer paths than NOP pings.

Another point of interest is that there are fewer processed routers than processed hosts, while in 2002 and in 2003 there were more routers than hosts. This difference leads to the conclusion that the average paths in the 2004 measurement were shorter than in the previous measurements. Figure 10 shows the different amounts of processed hosts in the three

Table 1: Measurements results

	2002	2003	2004
weighted average per-host-average NONE	**198.521639**	**145.230768**	**98.665882**
weighted average per-host-average NOP	**229.134040**	**165.180760**	**132.880307**
weighted average per-host-median NONE	**194.188480**	**145.159222**	**98.587564**
weighted average per-host-median NOP	**202.698017**	**147.551178**	**116.419992**
processed hosts and routers	**10153**	**9595**	**34896**
processed routers (with NOP-A)	**5726**	**5194**	**16850**
processed routers (without NOP-A)	**3226**	**3041**	**14508**
processed hosts only (NOP and NOP-A)	**4427**	**4401**	**18046**
NOP-A hosts	**1337**	**1490**	**5157**
hosts really interesting (NOP)	**3090**	**2907**	**12887**
pinged hosts *(with 10 pings measurement in 2002)*	**5759** *(27689)*	**27689**	**46326**
hosts with timeouts on NOP pings only	**306** *(4248)*	**3507**	**7545**

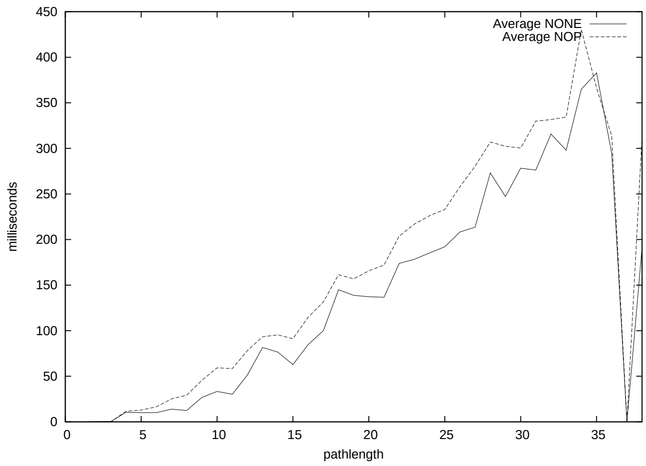

Figure 8: Averages per path length

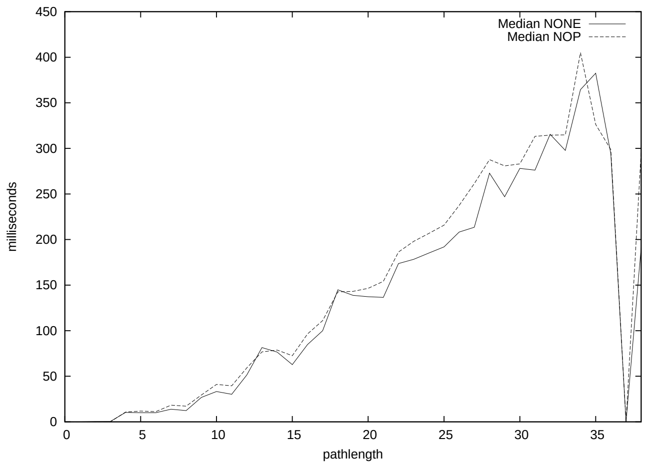

Figure 9: Medians per path length

measurements. The number of NONE hosts is the same as all answering hosts of the measurements because no host answered with IP options in response to NONE pings.

Fig. 11 shows the differences between the variances. The high variance of 2002 indicates problems at the provider in Linz for the mesurement of 2002. The variances of 2004 and 2003 are almost equal, even if in 2004 it is a bit higher. For NOP-A the variance is almost the double.

5. CONCLUSIONS

The measurement of 2004 was the last of a number of measurements, and the results did not vary significantly. The bigger and more current hostlist made some differences in terms of the quality of the measurement, but the statistics remained almost the same as in 2003 and 2002.

The difference between NONE and NOP pings of **26%** is a bit surprising, because the measurements of 2002 and 2003 show a value around **7-10%** — but the RTTs were generally lower in 2004, and lower values usually increase the relative difference. These measurements prove that NOP pings are processed slower than NONE, and they even show that some routers/hosts discard packets with IP options. They indicate that NOP pings are only up to **26%** slower than NONEs. This means that it could be still effective to enable certain options for special needs.

6. REFERENCES

[1] S. Ailleret. larbin.sourceforge.net.

[2] M. Rossi and M. Welzl. On the impact of ip option processing. University of Innsbruck, ICS - Institute of Computer Science, January 2004.

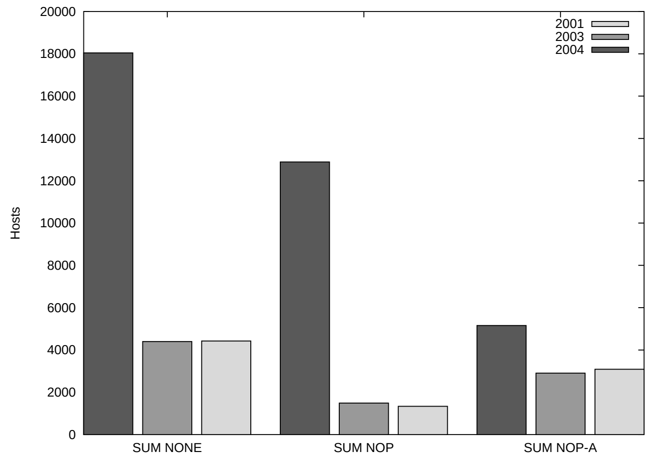

Figure 10: Difference of answering hosts per ping type between 2002, 2003 and 2004

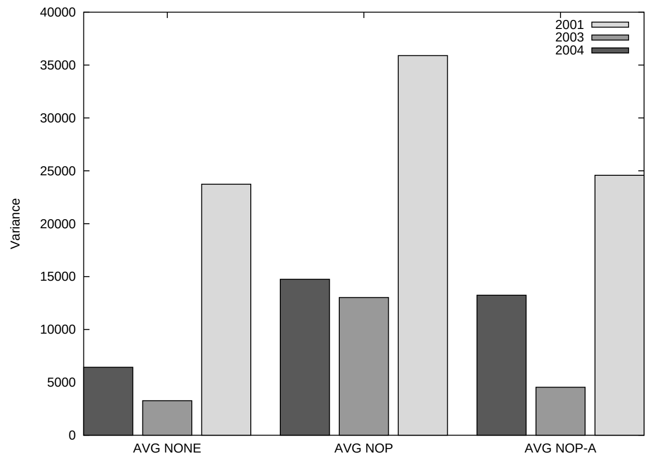

Figure 11: Variance differences between 2002, 2003 and 2004